arXiv:1109.6925v1 [cs.DC] 30 Sep 2011

Distributed Selfish Load Balancing with Weights and Speeds

Clemens Adolphs[1,*] **and Petra Berenbrink**[2]

1 **RWTH Aachen University, Aachen, Germany**
2 **Simon Fraser University, Burnaby, Canada**

Subject: Distributed Algorithms

Abstract

In this paper we consider neighborhood load balancing in the context of selfish clients. We assume that a network of n processors and m tasks is given. The processors may have different speeds and the tasks may have different weights. Every task is controlled by a selfish user. The objective of the user is to allocate his/her task to a processor with minimum load.

We revisit the concurrent probabilistic protocol introduced in [6], which works in sequential rounds. In each round every task is allowed to query the load of one randomly chosen neighboring processor. If that load is smaller the task will migrate to that processor with a suitably chosen probability. Using techniques from spectral graph theory we obtain upper bounds on the expected convergence time towards approximate and exact Nash equilibria that are significantly better than the previous results in [6]. We show results for uniform tasks on non-uniform processors and the general case where the tasks have different weights and the machines have speeds. To the best of our knowledge, these are the first results for this general setting.

Keywords and phrases Load balancing, reallocation, equilibrium, convergence

1 Introduction

Load Balancing is an important aspect of massively parallel computations as it must be ensured that resources are used to their full efficiency. Quite often the major constraint on balancing schemes for large networks is the requirement of *locality* in the sense that processors have to decide if and how to balance their load with local load information only. Global information is often unavailable and global coordination usually very expensive and impractical. Protocols for load balancing should respect this locality and still guarantee fast convergence to balanced states where every processor has more or less the same load.

In this paper we consider neighborhood load balancing in a selfish setting. We assume that a network of n processors and m tasks is given. The processors can have different speeds and the tasks can have different weights. Initially, each processor stores some number of tasks. The total number of tokens is time-invariant, i.e., neither do new tokens appear, nor do existing ones disappear. The *load* of a node at time t is the total weight of all tasks assigned to that node at that time.

Every task is assumed to belong to a selfish user. The goal of the user is to allocate the task to a processor with minimum load. We assume neighborhood load balancing, meaning that task movements are restricted by the network. Users that are assigned to the processor represented by node v of the network are only allowed to migrate their tasks over to processors

---

[*] The research was carried out during a visit to SFU

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

that are represented by neighboring nodes of v. Hence, the network models load balancing restrictions. Our model can be regarded as the selfish version of diffusion load balancing.

In this paper we revisit the concurrent probabilistic protocol introduced in [6]. The load balancing process works in sequential rounds. In each round every task is allowed to check the load of one randomly chosen neighboring processor. If that load is smaller the task will migrate to that processor with a suitably chosen probability. Note that, if the probability is too large (for example all tasks move to a neighbor with smaller load) the system would never be able to reach a balanced state. Here, we chose the migration probability as a function of the load difference of the two processors. No global information is necessary.

Using techniques from spectral graph theory similar to those used in [11], we can calculate upper bounds on the expected convergence time towards approximate and exact Nash equilibria that are significantly better than the previous results in [6]. We show results for uniform tasks on non-uniform processors and the general case where the tasks have different weights and the machines have speeds. To our best knowledge these are the first results for this general setting. For weighted tasks we deviate from the protocol for weighted tasks given in [6]. In our protocol, a player will move from one node to another only if the player with the largest weight would also do so. It is also straightforward to apply our techniques to discrete diffusive load balancing where each node sends the rounded expected flow of the randomized protocol to its neighbors ([2]).

1.1 Model and New Results

The computing network is represented by an undirected graph $G = (V, E)$ with vertices representing the processors and edges representing the direct communication links between them. The number of processors $n = |V|$ and the number of tasks is m. The *degree* of a vertex $v \in V$ is $\deg(v)$. The maximum degree of the network is denoted by Δ, and for two nodes v and w the maximum of $\deg(v)$ and $\deg(w)$ is d_{vw}.

$s_i \in \mathbb{R}$ is the speed of processor i. We assume that the speeds are scaled so that the smallest speed, called $s_{\min}$, is 1. If all speeds are the same we say the speeds are uniform. Let $\mathcal{S} = \sum_{i \in V} s_i$ If all s_i be the total capacity of the processors. Define $s_{\max}$ as the maximum speed and $s_{\min}$ as the minimum speed of the processors. In the case of weighted task task ℓ has a weight $w_l \in (0, 1]$. In the case of uniform tasks we assume the weight of all tasks is one. Let W denote the total sum of all weights, $W = \sum_i W_i(x)$.

A *state* x of the system is defined by the distribution of tasks among the processors. For the case of uniform indivisible tasks, we denote with $w_i(x)$ the number of tasks on processor i in state x. For the case of weighted tasks, $W_i(x)$ denotes the total weight on processor i whereas $w_l \in (0, 1]$ denotes the weight of tasks l. The *load* of a processor i is defined as $w_i(x)/s_i$ in the case of uniform tasks and as $W_i(x)/s_i$ in the case of weighted tasks. The goal is to reach a state x in which no task can benefit from migrating to a neighboring processor. Such s state is called *Nash Equilibrium*.

1.1.1 Uniform Tasks on Machines with Speeds

For uniform tasks, one round of the protocol goes as follows. Every task selects a neighboring node uniformly at random. If migrating to that node would lower the load experienced by the task, the task migrates to that node with proportional to the load difference and the speeds of the processors. For a detailed description of the protocol see Algorithm 1 in Section 3.

The first result concerns convergence to an approximate Nash equilibrium if the number of tasks, m, is large enough. For a detailed definition of Laplacian matrix see Section 2.

▶ **Theorem 1.1.** *Let* $\psi_c = 16n \cdot \Delta \cdot s_{\max}/\lambda_2$ *and let* λ_2 *denote the second smallest eigenvalue of the network's Laplacian matrix. Then, Algorithm 1 (p. 5) reaches a state* x *with* $\Psi_0(x) \leq 4 \cdot \psi_c$ *in expected time*

$$\mathcal{O}\left(\ln\left(\frac{m}{n}\right) \cdot \frac{\Delta}{\lambda_2} \cdot s_{\max}^2\right).$$

If $m \geq 8 \cdot \delta \cdot s_{\max} \cdot \mathcal{S} \cdot n^2$ *for some* $\delta > 1$*, this state is an* ε*-approximate-Nash equilibrium with* $\varepsilon = 2/(1+\delta)$*.*

From the state reached in Theorem 1.1, we then go on to prove the following bound for convergence to a Nash equilibrium.

▶ **Theorem 1.2.** *Let* ψ_c *be defined as in Theorem 1.1, and let* T *be the first time step in which the system is in a Nash equilibrium. Under the condition that the speeds* s_i *are integer multiples of a common factor,* ϵ*, it holds*

$$\mathbf{E}[T] = \mathcal{O}\left(n \cdot \frac{\Delta^2}{\lambda_2} \cdot \frac{s_{\max}^4}{\epsilon^2}\right).$$

These theorems are proven in Section 3. Our bound of Theorem 1.2 is smaller by at least a factor of $\Omega(\Delta \cdot \mathrm{diam}(G))$ than the bound found in [6] (see Observation 3.28).

Table 1 Comparison with existing results

Graph	ε-approximate NE		Nash Equilibrium	
	This Paper	[6]	This Paper	[6]
Complete Graph	$\ln\left(\frac{m}{n}\right)$	$n^2 \cdot \ln(m)$	n^2	n^6
Ring, Path	$n^2 \cdot \ln\left(\frac{m}{n}\right)$	$n^3 \cdot \ln(m)$	n^3	n^5
Mesh, Torus	$n \cdot \ln\left(\frac{m}{n}\right)$	$n^2 \cdot \ln(m)$	n^2	n^4
Hypercube	$\ln(n) \cdot \ln\left(\frac{m}{n}\right)$	$n \cdot \ln^3(n) \cdot \ln(m)$	$n \cdot \ln^2(n)$	$n^3 \cdot \ln^5(n)$

We summarize the results for the most important graph classes in Table 1. The table gives an overview of asymptotic bounds on the expected runtime to reach an approximate or a exact Nash equilibrium. We omit the speeds from this table because they are independent of the graph structure and, therefore, the same for each column. We compare the results of this paper to the bounds obtained from [6]. These contain a factor $\mathcal{S} = \sum_i s_i$, which we replace with n, using $\mathcal{S} = \sum_i s_i \geq n$. The table shows that for the graph classes at hand, our new bounds are superior to those in [6].

1.1.2 Weighted Tasks on Machines with Speeds

In Section 4, we study a slightly modified protocol (see 2) that allows tasks only to migrate to a neighboring processor if that would decrease their experienced load by a threshold depending on the speed of the processors. This protocol allows the tasks only to reach an approximate Nash Equilibrium.

▶ **Theorem 1.3.** *Let* $\psi_c = 16 \cdot n \cdot \Delta/\lambda_2 \cdot s_{\max}/s_{\min}^2$ *and let* λ_2 *denote the second smallest eigenvalue of the network's Laplacian matrix. Then, Algorithm 2 (p. 11) reaches a state* x *with* $\Psi_0(x) \leq 4 \cdot \psi_c$ *in time*

$$\mathcal{O}\left(\ln\left(\frac{m}{n}\right) \cdot \frac{\Delta}{\lambda_2} \cdot \frac{s_{\max}^2}{s_{\min}}\right).$$

Under the condition that $W > 8 \cdot \delta \cdot s_{\max}/s_{\min} \cdot \mathcal{S} \cdot n^2$ *for some* $\delta > 1$*, this state is an* $2/(1+\delta)$*-approximate Nash equilibrium.*

For the case of uniform speeds the theorem gives a bound of $\mathcal{O}\left(\ln(m/n) \cdot \Delta/\lambda_2\right)$ for the convergence time.

1.1.2.1 Outline.

After presenting the notation and preliminaries in Section 2, we treat the case of machines with speeds in Section 3. Section 4 treats the case of weighted tasks. Proofs are found in the appendix.

1.2 Related Work

The work closest to ours is in [4, 5, 6]. [4] considers the case of identical machines in a complete graph. The authors introduce a protocol similar to ours that reaches a Nash Equilibria (NE) in time $\mathcal{O}(\log\log m + \mathrm{poly}(n))$. Note that for complete graphs the NE and the optima (where the load discrepancy is zero or one) are identical. An extension of this model to weighted tasks is studied in [5]. Their protocol converges to a NE in time polynomially in n, m, and the largest weight. In [6] the authors consider a model similar to ours, meaning general graphs with processors with speed and weighted tasks. They use a potential function similar to ours for the analysis. The potential drop is linked to the maximum load deviation, $L_\Delta = \max_{i\in V} |e_i/s_i|$. The authors show that an edge must exist over which the load difference is at least $L_\Delta/\operatorname{diam}(G)$. As long as the potential Ψ_0 is large enough, it can then be shown that there is a multiplicative drop. This is then used to prove convergence to an approximate Nash equilibrium. Subsequently, a constant drop in Φ_1 is used to finally converge to a Nash equilibrium. The two main results of [6] for machines with speeds are presented in Table 1

Our paper relates to a general stream of works for selfish load balancing on a complete graph. There is a variety of issues that have been considered, starting with seminal papers on algorithms and dynamics to reach NE [13, 15]. More directly related are concurrent protocols for selfish load balancing in different contexts that allow convergence results similar to ours. Whereas some papers consider protocols that use some form of global information [14] or coordinated migration [19], others consider infinitesimal or splittable tasks [18, 3] or work without rationality assumptions [17, 1]. The machine models in these cases range from identical, uniformly related (linear with speeds) to unrelated machines. The latter also contains the case when there are access restrictions of certain agents to certain machines. For an overview of work on selfish load balancing see, e.g., [27].

Our protocol is also related to a vast amount of literature on (non-selfish) load balancing over networks, where results usually concern the case of identical machines and unweighted tasks. In expectation, our protocols mimic continuous diffusion, which has been studied initially in [10, 8] and later, e.g., in [25]. This work established the connection between convergence, discrepancy, and eigenvalues of graph matrices. Closer to our paper are discrete diffusion processes – prominently studied in [26], where the authors introduce a general technique to bound the load deviations between an idealized and the actual processes. Recently, randomized extensions of the algorithm in [26] have been considered, e.g., [12, 20].

2 Notation and Preliminaries

In this section we will give the more technical definitions.

A *state* x of the system is defined by the distribution of tasks among the processors. For the case of uniform indivisible tasks, we denote with $w_i(x)$ the number of tasks on processor i in state x. For the case of weighted tasks, $W_i(x)$ denotes the total weight on processor i whereas $w_l \in (0,1]$ denotes the weight of tasks l.

The *task vector* is defined as $\mathbf{w}(x) = (w_1(x), w_2(x), \cdots w_n(x))^\top$. We define the *load* of processor i in state x as $\ell_i(x) := w_i(x)/s_i$. In analogy to the task vector, we define the *load vector* as $\boldsymbol{\ell}(x) = (\ell_1(x), \cdots \ell_n(x))^\top$. For the processor speeds, we define the *speed vector* as $\mathbf{s} = (s_1, \cdots s_n)^\top$ and the *speed matrix* as $S \in \mathbb{N}^{n \times n}, \quad S_{ii} = s_i$. Let $s_{\max} := \max_{i \in V} s_i$ denote the maximum speed. The task vector $\mathbf{w}(x)$ and the load vector $\boldsymbol{\ell}(x)$ are related by the speed matrix S via $\boldsymbol{\ell}(x) = S^{-1}\mathbf{w}(x)$. The average load of the network is $\bar{\ell}_i = m/\mathcal{S}$. In the completely balanced state, each node has exactly this load. The corresponding task vector is $\bar{\mathbf{w}} = m/\mathcal{S} \cdot \mathbf{s}$ and we define $\mathbf{e}(x)$ of the deviation of the actual task vector from the average load vector, $\mathbf{e}(x) = \mathbf{w}(x) - \bar{\mathbf{w}}$. It is clear that $\sum_{i \in V} e_i = 0$.

A state x of the system is called a *Nash equilibrium* (NE) if no single task can improve its perceived load by migrating to a neighboring node while all other tasks remain where they are, i.e., $\ell_i - \ell_j \le 1/s_j$ for all edges (i,j). A state x of the system is called an ε*-approximate Nash equilibrium* (ε-approximate-NE) if no single task can improve its perceived load by a factor of $(1-\varepsilon)$, i.e. $(1-\varepsilon)\cdot \ell_i - \ell_j \le 1/s_j$.

The Laplacian $L(G)$ is a matrix widely used in graph theory. It is the $n \times n$ matrix whose diagonal elements are $L_{ii} = \deg(i)$, and the off-diagonal elements are $L_{ij} = -1$ if $(i,j) \in E(G)$ and 0 otherwise. The generalized Laplacian LS^{-1}, where S is the diagonal matrix containing the speeds s_i [11], is used to analyze the behavior of migration in heterogeneous networks.

3 Uniform Tasks on Machines with Speeds

The pseudo-code of our protocol is given in Algorithm 1. Recall that $d_{i,j}$ is defined as $\max\{\deg(i), \deg(j)\}$. α is defined as $4s_{\max}$.

Algorithm 1: Distributed Selfish Load Balancing

```
begin
    foreach task $\ell$ in parallel do
        Let $i = i(l)$ be the current machine of task $l$
        Choose a neighboring machine $j$ uniformly at random
        if $\ell_i - \ell_j > 1/s_j$ then
            Move task $\ell$ from node $i$ to node $j$ with probability
                $p_{ij} := \frac{\deg(i)}{d_{i,j}} \cdot \frac{\ell_i - \ell_j}{\alpha \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot W_i}$
        end
    end
end
```

The analysis of this protocol initially follows the steps of [6] up to Lemma 3.3 (Restated as Lemma 4.1 in the appendix). Before we outline the remainder of our proof, we introduce some more notation.

▶ **Definition 3.1.** For a given state x, we define $f_{ij}(x)$ as the *expected flow* over edge (i, j). It holds

$$f_{ij}(x) = \begin{cases} \dfrac{\ell_i(x) - \ell_j(x)}{\alpha \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right)} & \text{if} \quad \ell_i(x) - \ell_j(x) > \frac{1}{s_j} \\ 0 & \text{otherwise.} \end{cases}$$

The following two potential functions will be used in the analysis.

▶ **Definition 3.2.** For $r = 0, 1$, define

$$\Phi_r(x) := \sum_{i \in V} \frac{W_i(x) \cdot (W_i(x) + r)}{s_i}.$$

The potential Φ_0 is minimized for the average task vector, $\bar{\mathbf{w}}$. We define the according normalized potential Ψ_0.

▶ **Definition 3.3.** The normalized potential $\Psi_0(x)$ is defined as

$$\Psi_0(x) = \Phi_0(x) - \frac{m^2}{\mathcal{S}} = \sum_{i \in V} \frac{e_i(x)^2}{s_i}.$$

We want to relate this potential function to the load imbalance in the system. To this end, we define a new quantity.

▶ **Definition 3.4.** We define the *maximum load difference* as

$$L_\Delta(x) = \max_{i \in V} \left| \frac{W_i(x)}{s_i} - \frac{m}{\mathcal{S}} \right| = \max_{i \in V} \left| \frac{e_i}{s_i} \right|.$$

▶ **Definition 3.5.** Let $t > 0$ be some time step during the executing of our protocol and let X^t denote the state of the system at that time step. We define $\Delta\Phi_r(X^t) := \Phi_r(X^{t-1}) - \Phi_r(X^t)$ as the drop in potential $\Phi_r(X^t)$ in time step t. The sign convention for $\Delta\Phi_r(X^t)$ is such that a *drop* in $\Phi_r(x)$ from time step $t-1$ to t gets a *positive* sign. This emphasizes that a large drop in $\Phi_r(x)$ is a desirable outcome of our process. $\Delta\Psi_0(X^t)$ is defined analogously.

▶ **Lemma 3.6.** *The shifted potential $\Psi_0(x)$ has the following properties.*

(1) The change in $\Psi_0(x)$ due to migrating tasks is the same as the change in $\Phi_0(x)$, i.e.

$$\Delta\Psi_0(X^t | X^{t-1} = x) = \Delta\Phi_0(X^t | X^{t-1} = x)$$

(2) The potential $\Psi_0(x)$ can also be written using the generalized dot-product introduced in Section A.2, $\Psi_0(x) = \sum_{i \in V} e_i^2 / s_i = \langle \mathbf{e}, \mathbf{e} \rangle_S$

▶ **Definition 3.7.** With $f_{ij}(x)$ the expected flow over edge (i, j) in state x, we define the set of *non-Nash edges* as

$$\tilde{E}(x) := \{(i, j) \in E : f_{ij}(x) > 0\}.$$

This is the set of edges for which tasks have an incentive to migrate. Edges with $f_{ij}(x) = 0$ are called *Nash edges* or *balanced edges.*

▶ **Definition 3.8.** As an auxiliary quantity, we define

$$\Lambda^r_{ij}(x) := (2\alpha - 2) \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot f_{ij}(x) + \frac{r}{s_i} - \frac{r}{s_j}.$$

Our improved bound builds upon results in [6]. In that paper, the randomized process is analyzed by first lower-bounding the potential drop in the case that exactly the expected number of tasks is moved, and then by upper-bounding the variance of that process. This leads to Lemma 4.1. Based on this lemma, we now prove a stronger bound on the expected drop in the potential $\Psi_0(x)$. Let us briefly outline the necessary steps. The lower bound on the drop in the potential in Lemma 4.1 is a sum over the non-Nash edges and contains terms of the form $\ell_i - \ell_j$, whereas the potential itself is a sum over the nodes and contains terms of the form ℓ_i^2. We will use the graph's Laplacian matrix to establish a connection between Ψ_0 and the expected drop in Ψ_0. This will allow us to prove fast convergence to a state where Ψ_0 is below a certain critical value ψ_c. If m is sufficiently large, this state also is an ε-approximate Nash equilibrium. In the next stage of our approach, we use a constant drop in Ψ_1, a shifted version of Φ_1, to prove convergence to an exact Nash equilibrium. The techniques from probability theory used in this this section are similar to the ones used in [4].

3.1 Convergence Towards an Approximate Nash Equilibrium

To make the connection with the Laplacian, we first have to rewrite the bound in Lemma 4.1 in the following way.

▶ **Lemma 3.9.** *Under the condition that the system is in state x, the expected drop in the potentials Φ_0 and Ψ_0 is bounded by*

$$\mathbf{E}[\Delta\Psi_0(X^{k+1})|X^k = x] \geq \sum_{(i,j)\in E} \left[\frac{(1 - \frac{2}{\alpha}) \cdot (\ell_i(x) - \ell_j(x))^2}{\alpha \cdot d_{i,j} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right)}\right] - \frac{n}{\alpha}.$$

Next, we use various technical results from spectral graph theory to prove the following bound.

▶ **Lemma 3.10.** *Let L be the Laplacian of the network. Let λ_2 be its second smallest eigenvalue. Then*

$$\mathbf{E}[\Delta\Psi_0(X^{k+1})|X^k = x] \geq \frac{\lambda_2}{16\Delta} \cdot \frac{1}{s^2_{\max}} \cdot \Psi_0 - \frac{n}{4 \cdot s_{\max}}.$$

In a first step, we get rid of the conditioning of the potential drop on the previous state.

▶ **Lemma 3.11.** *Let γ be defined such that $1/\gamma = \lambda_2/(32\Delta \cdot s^2_{\max})$.*
Then, the expected value of the potential in time step t is at most

$$\mathbf{E}[\Psi_0(X^t)] \leq \left(1 - \frac{2}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^{t-1})] + \frac{n}{4 \cdot s_{\max}}.$$

As long as the expected value of the potential is sufficiently large, we can rewrite the potential drop as a multiplicative drop.

▶ **Definition 3.12.** Let λ_2 be the second smallest eigenvalue of the Laplacian $L(G)$ of the network. We define the *critical value* ψ_c as $\psi_c = 8 \cdot n \cdot \Delta \cdot s_{\max}/\lambda_2$.

▶ **Lemma 3.13.** *Let t be a time step for which the expected value of the potential satisfies $\mathbf{E}[\Psi_0(X^t)] \geq \psi_c$. Let γ be defined as in Lemma 3.11. Then, the expected potential in time step $t+1$ is bounded by*

$$\mathbf{E}[\Psi_0(X^{t+1})] \leq \left(1 - \frac{1}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^t)].$$

This immediately allows us to prove the following.

▶ **Lemma 3.14.** *For a given time step T, there either is a $t < T$ so that $\mathbf{E}[\Psi_0(X^t)] \leq \psi_c$, or*

$$\mathbf{E}[\Psi_0(X^T)] \leq \left(1 - \frac{1}{\gamma}\right)^T \cdot \mathbf{E}[\Psi_0(X^0)].$$

Thus, as long as $\mathbf{E}[\Psi_0(X^t)] > \psi_c$ holds, the expected potential drops by a constant factor. This allows us to derive a bound on the time it takes until $\mathbf{E}[\Psi_0(X^t)]$ is small.

▶ **Lemma 3.15.** *Let $T = 2\gamma \cdot \ln(m/n)$. Then it holds*

(1) There is a $t \leq T$ such that $\mathbf{E}[\Psi_0(X^T)] \leq \psi_c$.

(2) There is a $t \leq T$ such that the probability that $\Psi_0(X^t) \leq 4 \cdot \psi_c$ is at least

$$\mathbf{Pr}[\Psi_0(X^t) \leq 4 \cdot \psi_c] \geq \frac{3}{4}.$$

This is similar to a result in [6], but our factor γ is different. This is reflected in a different expected time needed to reach an ε-approximate Nash equilibrium, as we have pointed out in the introduction.

Next, we show that states with $\Psi_0(x) \leq 4 \cdot \psi_c$ are indeed ε-approximate Nash equilibria if the number of tasks exceeds a certain threshold. This requires one further observation.

▶ Observation 3.16. For any state x, we have $L_\Delta(x)^2 \leq \Psi_0(x) \leq \mathcal{S} \cdot L_\Delta(x)^2$.

▶ **Lemma 3.17.** *Let $m \geq 8 \cdot \delta \cdot n^2 \cdot \mathcal{S} \cdot s_{\max}$ for some $\delta > 1$. Then a state x with $\Psi_0(x) \leq 4 \cdot \psi_c$ is a $2/(1+\delta)$-approximate Nash equilibrium.*

▶ Remark. If m is small, it still holds that we reach a state x with $\Psi_0(x) \leq 4 \cdot \psi_c$, which is all we need to prove convergence to an exact Nash equilibrium in the next section. It is just that this intermediate state is then not an ε-approximate-Nash equilibrium.

Now we are ready to show Theorem 1.1.

Theorem 1.1. Lemma 3.17 ensures that after T steps the probability for *not* having reached a state x with $\Psi_0(x) \leq 4 \cdot \psi_c$ is at most 1/4. Hence, the expected number of times we have to repeat T steps is less than

$$1 + 1/4 + 1/4^2 + \cdots = \frac{1}{1 - \frac{1}{4}} < 2.$$

The expected time needed to reach such a state is therefore at most $2 \cdot T$ with T from Lemma 3.15. ◀

If we let the algorithm iterate until a state x with $\Psi_0(x) \leq 4 \cdot \psi_c$ is obtained, Theorem 1.1 bounds the *expected* number of time steps we have to perform. However, by repeating a sufficient number of blocks with T steps, we can obtain arbitrary high probability.

▶ **Corollary 3.18.** *After $c \cdot \log_4 n$ many blocks of size T, a state with $\Psi_0(x) \leq 4 \cdot \psi_c$ is reached with probability at least $1 - 1/n^c$.*

Corollary 3.18. The probability for *not* reaching a state x with $\Psi_0(x) \leq 4 \cdot \psi_c$ after k steps is at most $1/4^k$. We are interested in the complementary event, so its probability is at least $1 - 1/4^k$. For $k = c \cdot \log_4 n$ the statement follows immediately. ◀

3.2 Convergence Towards a Nash Equilibrium

We now prove the upper bound for the expected time necessary to reach an exact Nash Equilibrium (Theorem 1.2, p. 3). To show this result, we have to impose a certain condition on the speeds. If the speeds are arbitrary non-integers, convergence can become arbitrarily slow. Therefore, we assume that there exists a common factor $\epsilon \in (0, 1]$ so that for every speed s_i there exists an integer $n_i \in \mathbb{N}$ so that $s_i = n_i \cdot \epsilon$. We call ϵ the *granularity* of the speed distribution. The convergence factor α, which was $4s_{\max}$ in the original protocol, must be changed to $4s_{\max}/\epsilon$. For non-integer speeds, we have $\epsilon < 1$, so this effectively increases α.

To show convergence towards an exact Nash Equilibrium we cannot rely solely on the potential $\Psi_0(x)$, because when the system is close to a Nash equilibrium it is possible that the potential function increases even when a task makes a move that improves its perceived load. Therefore, we now look at potential $\Phi_1(x)$.

▶ **Definition 3.19.** We define the *shifted potential function*

$$\Psi_1(x) = \Phi_1(x) - \frac{m^2}{\mathcal{S}} - \frac{m \cdot n}{\mathcal{S}} - \frac{n^2}{4\mathcal{S}} + \frac{1}{4}\sum_i \frac{1}{s_i}.$$

Let $\bar{s}_a$ and $\bar{s}_h$ denote the arithmetic mean and the harmonic mean of the speeds, i.e., $s_a = \sum_{i \in V} s_i/n$ and $s_h = n/\sum_{i \in V} 1/s_i$.

Then, we can write

$$\Psi_1(x) = \Phi_1(x) - \frac{m^2}{\mathcal{S}} - \frac{m \cdot n}{\mathcal{S}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right).$$

▶ Observation 3.20. The shifted potential $\Psi_1(x)$ has the following properties.

(1) Let $\mathbf{e} = \mathbf{w} - \bar{\mathbf{w}}$ be the task deviation vector. Then

$$\Psi_1(x) = \sum_{i \in V} \left[\frac{\left(e_i + \frac{1}{2}\right)^2}{s_i}\right] - \frac{n}{4\bar{s}_a}.$$

(2) $\Psi_1(x) \geq 0$.

(3) $\Psi_1(x) = \Psi_0(x) + \sum_{i \in V} \frac{e_i}{s_i} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right)$.

(4) $\Delta\Psi_1(X^t) = \Delta\Phi_1(X^t)$.

Before we can lower-bound the expected drop in $\Psi_1(x)$, we need a technical lemma regarding a lower bound to the load difference. It is similar to [6, Lemma 3.7], which concerned integer speeds, so the result here is more general.

▶ **Lemma 3.21.** *Every edge (i, j) with $\ell_i - \ell_j > 1/s_j$ also satisfies*

$$\ell_i - \ell_j \geq \frac{1}{s_j} + \frac{\epsilon}{s_i \cdot s_j}.$$

Potential Ψ_1 differs from potential Ψ' defined in [6] by a constant only. Therefore, potential differences are the same for both potentials and we can apply results for Ψ' to Ψ_1.

▶ **Lemma 3.22.** *If the system is in a state x that is not a Nash equilibrium, then*

$$\mathbf{E}[\Delta\Psi_1(X^{k+1}) | X^k = x] \geq \frac{\epsilon^2}{8\Delta \cdot s_{\max}^3}$$

Since the results of the previous section apply to Ψ_0 whereas now we work with Ψ_1, we add this technical lemma relating the two.

▶ **Lemma 3.23.** *For any state x it holds*

$$\Psi_1(x) \leq \Psi_0(x) + \sqrt{\Psi_0(x) \cdot \frac{n}{\bar{s}_h}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right).$$

To obtain a bound on the expected time the system needs to reach the NE, we use a standard argument from martingale theory. Let us abbreviate $V := \epsilon^2/(8\Delta \cdot s_{\max}^3)$. We introduce a new random variable Z_t which we define as $Z_t = \Psi_1(X^t) + t \cdot V$.

▶ **Lemma 3.24.** *Let T be the first time step for which the system is in a Nash equilibrium. Then, for all times $t \leq T$ we have*
(1) $\mathbf{E}[Z_t | Z_{t-1} = z] \leq z$
(2) $\mathbf{E}[Z_t] \leq \mathbf{E}[Z_{t-1}]$.

▶ **Corollary 3.25.** *Let T be the first time step for which the system is in a Nash equilibrium. Let $t \wedge T$ be defined as $\min\{t, T\}$. Then the random variable $Z_{t\wedge T}$ is a super-martingale.*

▶ **Corollary 3.26.** *Let T be the first time step for which the system is in a Nash equilibrium. Then $\mathbf{E}[Z_T] \leq Z_0 = \Psi_1(X^0)$.*

Now we are ready to show Theorem 1.2.

Theorem 1.2. First, we assume that at time $t = 0$ the system is in a state with $\mathbf{E}[\Psi_0(X^T)] \leq 4 \cdot \psi_c$. Using the non-negativity of $\Psi_1(x)$ (Observation 3.20) allows us to state

$$\begin{aligned}
V \cdot \mathbf{E}[T] &\leq \mathbf{E}[\Psi_1(X^T)] + V \cdot \mathbf{E}[T] = \mathbf{E}[Z_T] \\
\text{(Cor. 3.26)} \qquad &\leq \mathbf{E}[Z_0] = \Psi_1(X^0) \\
\text{(Lem. 3.23)} \qquad &\leq \Psi_0(X^0) + \sqrt{\Psi_0(X^0) \cdot \frac{n}{\bar{s}_h}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right) \\
&\leq 4 \cdot \psi_c + \sqrt{4 \cdot \psi_c \cdot \frac{n}{\bar{s}_h}} + \frac{n}{4}
\end{aligned}$$

Inserting the definition of ψ_c and dividing by V yields

$$\begin{aligned}
\mathbf{E}[T] &\leq 8\Delta \cdot \frac{s_{\max}^3}{\epsilon^2} \cdot \left[4 \cdot \frac{16 \cdot n \cdot \Delta \cdot s_{\max}}{\lambda_2} + \sqrt{4 \cdot \frac{16 \cdot n \cdot \Delta \cdot s_{\max}}{\lambda_2} \cdot n} + \frac{n}{4}\right] \\
\text{(Lem. 1.7)} \qquad &\leq 8\Delta \cdot \frac{s_{\max}^3}{\epsilon^2} \cdot \left[\frac{64 \cdot n \cdot \Delta \cdot s_{\max}}{\lambda_2} + \sqrt{32 \cdot n^2 \cdot s_{\max} \cdot} \cdot \frac{2\Delta}{\lambda_2} + \frac{n}{4}\right] \\
&\leq 512 \cdot \Delta^2 \cdot \frac{s_{\max}^4}{\epsilon^2} \cdot \frac{n}{\lambda_2} + 91 \cdot \Delta^2 \cdot \frac{s_{\max}^4}{\epsilon^2} \cdot \frac{n}{\lambda_2} + 4 \cdot \Delta^2 \cdot \frac{s_{\max}^4}{\epsilon^2} \cdot \frac{n}{\lambda_2} \\
&= 607 \cdot \Delta^2 \cdot \frac{s_{\max}^4}{\epsilon^2} \cdot \frac{n}{\lambda_2}.
\end{aligned}$$

where we have used that $2\Delta/\lambda_2 \geq 1$ (Lemma 1.7) to pull that expression outside of the square root in the first line.

This bound was derived under the assumption that at $t = 0$ we had a state with $\mathbf{E}[\Psi_0(X^t)] =\leq 4 \cdot \psi_c$. If this is not the case, let τ denote the number of time steps to reach such a state, and let T' denote the additional number of time steps to reach a NE from there. Combining the result from above with Theorem 1.1 allows us to write

$$\mathbf{E}[T] = \mathbf{E}[\tau + T'] = \mathcal{O}\left(\frac{n}{\lambda_2} \cdot \Delta^2 \cdot \frac{s_{\max}^4}{\epsilon^2}\right).$$

◀

▶ **Corollary 3.27.** *Similarly to Corollary 3.18, after $c \cdot \log_4 n$ blocks of T steps we have reached a Nash Equilibrium with probability at least $1 - 1/n^c$.*

▶ Observation 3.28. Our bound in Theorem 1.2 is asymptotically lower than the corresponding bound in [6] by at least a factor of $\Omega\left(\Delta \cdot \mathrm{diam}(G)\right)$.

Proof. Lemma 1.5 yields $n \cdot \mathrm{diam}(G) \geq 4/\lambda_2$. Additionally, we have $\mathcal{S} \geq s_{\max}$, since $s_{\max}$ occurs (at least once) in the sum of all speeds. Hence, the asymptotic bound from [6] is larger than

$$\mathcal{O}\left(n \cdot \frac{\Delta^2}{\lambda_2} \cdot s_{\max}^4 \cdot [\Delta \cdot \mathrm{diam}(G)]\right).$$

The first part of this expression is the bound of Theorem 1.2, so the expression in the square brackets is the additional factor of the bound from [6]. ◀

4 Weighted Tasks

The set of tasks assigned to node i is called $x(i)$. The weight of node i becomes $W_i(x) = \sum_{\ell \in x(i)} w_\ell$ whereas the corresponding load is defined as $\ell_i(x) = W_i(x)/s_i$.

We present a protocol for weighted tasks that differs from the one described in [6]. It is presented in Algorithm 2

Algorithm 2: Distributed Selfish Load Balancing for weighted tasks

```
begin
    foreach task $\ell$ in parallel do
        Let $i = i(l)$ be the current machine of task $\ell$
        Choose a neighboring machine $j$ uniformly at random
        if $\ell_i - \ell_j > \frac{1}{s_j}$ then
            Move task $\ell$ from node $i$ to node $j$ with probability
                $p_{ij} := \frac{\deg(i)}{d_{i,j}} \cdot \frac{W_i - W_j}{2\alpha \cdot W_i}$
        end
    end
end
```

The notable difference to the scheme in [6] is that in our case, the decision of a task ℓ to migrate or not does not depend on that task's weight. In the original protocol, a load difference of more than w_ℓ/s_j would suffice for task ℓ to have an incentive to migrate. In the modified protocol, a task will only move if the load difference is at least $1/s_j$. The advantage of this approach is that for an edge (i, j), either all or none of the tasks on node i have an incentive to migrate. This greatly simplifies the analysis. We will show that the system rapidly converges to a state where $\ell_i - \ell_j \leq 1/s_j$ for all edges (i,j). Such a system is not necessarily a Nash equilibrium as $\ell_i - \ell_j$ might still be larger than the size of a given task w_ℓ. We will show, however, that such a state is an ε-approximate NE.

▶ **Definition 4.1.** In analogy to the unweighted case, we define the expected flow f_{ij} as the expected *weight* of the tasks migrating from i to j in state x. It is given by

$$f_{ij}(x) = \frac{\ell_i(x) - \ell_j(x)}{\alpha \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot W_i(x)} \cdot \sum_{\ell \in x(i)} w_\ell = \frac{\ell_i(x) - \ell_j(x)}{\alpha \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right)}.$$

The potentials Φ_0 and Φ_1 are defined analogously to the unweighted case. Here, we concentrate on Φ_0 alone. The average weight per node is W/n and the task deviation e_i is defined as $W_i - W/n$. We define $\Psi_0(x)$ in analogy to the unweighted case as the normalized version of Φ_0,

$$\Psi_0 = \Phi_0 - \frac{W^2}{\sum_i s_i} = \sum_{i \in V} \frac{e_i^2}{s_i}.$$

The auxiliary quantity $\Lambda_{ij}(x)$ is defined analogously to the unweighted case as

$$\Lambda_{ij}(x) = (2\alpha - 2) \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot f_{ij}(x).$$

4.1 Convergence Towards an Approximate Nash Equilibrium

In close analogy to [6, Lemma 3.1], we first bound the drop of the potential when the flow is exactly the expected flow.

▶ **Lemma 4.2.** *The drop in potential Φ_0 if the system is in state x and if the flow is exactly the expected flow is bounded by*

$$\tilde{\Delta}\Phi_0(X^{t+1} | X^t = x) \geq \sum_{(i,j) \in \tilde{E}} f_{ij} \cdot \Lambda_{ij}.$$

The proof is formally equivalent to the one in [6] and therefore omitted here. Next, we bound the variance of the process.

▶ **Lemma 4.3.** *The variances of the weights on the nodes are bounded via*

$$\sum_i \frac{\mathbf{Var}[W_i(X^t) | X^{t-1} = x]}{s_i} \leq \sum_{ij} f_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right)$$

This allows us to formulate a bound on the expected potential drop in analogy to [6, Lemma 3.3] by combining Lemma 4.2 and Lemma 4.3.

▶ **Lemma 4.4.** *The expected drop in potential Φ_0 if the system is in state x is at least*

$$E[\Delta\Phi_0(X^t) | X^{t-1} = x] \geq \sum_{(i,j) \in \tilde{E}} f_{ij}(x) \cdot (\Lambda_{ij}(x) - 2).$$

The proof is analogous to the corresponding lemma in [6].

Theorem 1.3. The rest of the proof is the same as the proof for the unweighted case. One may verify that, indeed, Lemma 3.9 and all subsequent results do not rely on the specific form of ℓ_i or the underlying nature of the tasks. Using the same eigenvalue techniques as in the unweighted case, this allows us to obtain a bound involving the second smallest eigenvalue of the graph's Laplacian matrix. Following the steps of the unweighted case allows us to prove the main result of this section. ◀

Acknowledgements

The authors thank Thomas Sauerwald for helpful discussions.

References

1 Heiner Ackermann, Simon Fischer, Martin Hoefer, and Marcel Schöngens. Distributed algorithms for qos load balancing. In *Proceedings of the twenty-first annual symposium on Parallelism in algorithms and architectures*, SPAA'09, pages 197–203, New York, NY, USA, 2009. ACM.

2 Clemens P. J. Adolphs and Petra Berenbrink. Improved Bounds for Discrete Diffusive Load Balancing. Manuscript, 2012.

3 Baruch Awerbuch, Yossi Azar, and Rohit Khandekar. Fast load balancing via bounded best response. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 314–322, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.

4 Petra Berenbrink, Tom Friedetzky, Leslie Ann Goldberg, Paul W. Goldberg, Zengjian Hu, and Russell Martin. Distributed Selfish Load Balancing. *SIAM Journal on Computing*, 37(4):1163, 2007.

5 Petra Berenbrink, Tom Friedetzky, Iman Hajirasouliha, and Zengjian Hu. Convergence to equilibria in distributed, selfish reallocation processes with weighted tasks. In Lars Arge, Michael Hoffmann, and Emo Welzl, editors, *Proceedings of the 15th Annual European Symposium on Algorithms (ESA 2007)*, volume 4698/2007 of *Lecture Notes in Computer Science*, pages 41–52. Springer, Springer, October 2007.

6 Petra Berenbrink, Martin Hoefer, and Thomas Sauerwald. Distributed selfish load balancing on networks. In *Proceedings of 22nd Symposium on Discrete Algorithms (SODA'11)*, pages 1487–1497, 2011.

7 Rajendra Bhatia. Linear Algebra to Quantum Cohomology: The Story of Alfred Horn's Inequalities. *The American Mathematical Monthly*, 108(4):289 – 318, 2001.

8 J. E. Boillat. Load balancing and Poisson equation in a graph. *Concurrency: Practice and Experience*, 2(4):289–313, December 1990.

9 Fan R. K. Chung. *Spectral graph theory.* AMS Bookstore, 1997.

10 G Cybenko. Dynamic load balancing for distributed memory multiprocessors. *Journal of Parallel and Distributed Computing*, 7(2):279–301, October 1989.

11 Robert Elsässer, Burkhard Monien, and Robert Preis. Diffusion Schemes for Load Balancing on Heterogeneous Networks. *Theory of Computing Systems*, 35(3):305–320, May 2002.

12 Robert Elsässer, Burkhard Monien, and Stefan Schamberger. Distributing unit size workload packages in heterogeneous networks. *Journal of Graph Algorithms and Applications*, 10(1):51–68, 2006.

13 Eyal Even-Dar, Alex Kesselman, and Yishay Mansour. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3):32–es, August 2007.

14 Eyal Even-Dar and Yishay Mansour. Fast convergence of selfish rerouting. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms (SODA'05)*, pages 772–781, 2005.

15 Rainer Feldmann, Martin Gairing, Thomas Lücking, Burkhard Monien, and Manuel Rode. Nashification and the coordination ratio for a selfish routing game. In Jos Baeten, Jan Lenstra, Joachim Parrow, and Gerhard Woeginger, editors, *Automata, Languages and Programming*, volume 2719 of *Lecture Notes in Computer Science*, pages 190–190. Springer Berlin / Heidelberg, 2003.

16 Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

17 S. Fischer, P. Mahonen, M. Schongens, and B. Vocking. Load balancing for dynamic spectrum assignment with local information for secondary users. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1 –9, oct. 2008.

18 Simon Fischer and Berthold Vöcking. Adaptive routing with stale information. In *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing (PODC'05)*, pages 276–283, New York, NY, USA, 2005. ACM.

19 D. Fotakis, A. Kaporis, and P. Spirakis. Atomic congestion games: Fast, myopic and concurrent. *Theory of Computing Systems*, 47:38–59, 2010. 10.1007/s00224-009-9198-2.

20 Tobias Friedrich and Thomas Sauerwald. Near-perfect load balancing by randomized rounding. In *Proceedings of the 41st annual ACM symposium on Symposium on theory of computing - STOC '09*, page 121, New York, New York, USA, May 2009. ACM Press.

21 A. A. Klyachko. Random walks on symmetric spaces and inequalities for matrix spectra. *Linear Algebra and its Applications*, 319(1-3):37–59, November 2000.

22 Bojan Mohar. Isoperimetric numbers of graphs. *Journal of Combinatorial Theory, Series B*, 47(3):274–291, December 1989.

23 Bojan Mohar. Eigenvalues, diameter, and mean distance in graphs. *Graphs and Combinatorics*, 7(1):53–64, March 1991.

24 Mohar, B. The Laplacian Spectrum of Graphs. In Y. Alavi, editor, *Graph theory, combinatorics, and applications*, volume 2, pages 871–898. Wiley, 1991.

25 S Muthukrishnan, B. Ghosh, and M.H. Schultz. First- and Second-Order Diffusive Methods for Rapid, Coarse, Distributed Load Balancing. *Theory of Computing Systems*, 31(4):331–354, July 1998.

26 Y. Rabani, A. Sinclair, and R. Wanka. Local divergence of Markov chains and the analysis of iterative load-balancing schemes. In *Proceedings 39th Annual Symposium on Foundations of Computer Science (FOCS'98)*, pages 694–703. IEEE Comput. Soc, 1998.

27 Berthold Vöcking. Selfish Load Balancing. In Noam Nisan, Eva Tardos, Tim Roughgarden, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 20. Cambridge University Press, 2007.

28 Hermann Weyl. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, December 1912.

A Spectral Graph Theory

In this appendix, we will briefly summarize some important theorems of spectral graph theory. For an excellent introduction, we recommend the book by Fan Chung [9]. Many important results are collected in an overview article by Mohar [24].

Results in this section are, unless indicated otherwise, taken from these sources. Let us begin by defining the matrix we are interested in.

▶ **Definition 1.1.** Let $G = (V, E)$ be an undirected graph with vertices $V = \{1, \dots n\}$ and edges E.

The *Laplacian* $L(G)$ of G is defined as

$$L(G) \in \mathbb{N}^{n \times n} \qquad L(G)_{ij} = \begin{cases} \deg(i) & i = j \\ -1 & (i,j) \in E \\ 0 & \text{otherwise.} \end{cases}$$

The following Lemma summarizes some basic properties of $\tilde{L}(G)$ and, therefore, also of $L(G)$. These properties are found in every introduction to spectral graph theory.

▶ **Lemma 1.2.** *Let $L(G)$ be the Laplacian of a graph G. For brevity, we omit the argument G in the following. Then, L satisfies the following.*

(1) For every vector $\mathbf{x} \in \mathbb{R}^n$ we have

$$\mathbf{x}^\top L\mathbf{x} = \sum_{i,j\in V} x_i \cdot L_{ij} \cdot x_j = \sum_{(i,j)\in E} c_{ij} \cdot (x_i - x_j)^2$$

(2) L is symmetric positive semi-definite, i.e., $L^\top = L$ and $\mathbf{x}^\top L\mathbf{x} \geq 0$ for every vector $\mathbf{x}$.

(3) Each column (row) of $\tilde{L}$ sums to 0.

A.1 Spectral Analysis

We now turn our attention to the spectrum of the Laplacian.

▶ **Definition 1.3.** Let $L(G)$ be the Laplacian of a graph G. Lemma 1.2 and the spectral theorem of linear algebra ensure that L has an orthogonal eigenbasis, i.e. there are n (not necessarily distinct) eigenvalues with n linearly independent eigenvectors which can be chosen to be mutually orthogonal.

We call the eigenvalues of $L(G)$ the *Laplacian spectrum* of G and write

$$\lambda(G) = (\lambda_1 \leq \lambda_2 \cdots \leq \lambda_n)$$

where the λ_i are the eigenvalues of $L(G)$.

The corresponding eigenvectors are denoted $\mathbf{v}_i$.

The Laplacian spectrum of G contains valuable information about G. Some very basic results are given in the next Lemma.

▶ **Lemma 1.4.** *Let G be a graph with Laplacian spectrum $\lambda(G)$. For a graph $G = (V, E)$ the following holds for both the unweighted and the weighted spectrum.*

(1) The vector $\mathbf{1} := (1, \cdots, 1)^\top$ is eigenvector to L and $\tilde{L}$ with eigenvalue 0. Hence, $\lambda_1 = 0$ is always the smallest eigenvalue of any Laplacian.

(2) The multiplicity of the eigenvalue 0 is equal to the number of connected components of G. In particular, a connected graph has $\lambda_1 = 0$ and $\lambda_2 > 0$.

The second-smallest eigenvalue λ_2 is closely related to the connectivity properties of G. It was therefore called *algebraic connectivity* when it was first intensely studied by Fiedler [16]. The eigenvector corresponding to λ_2 is also called *Fiedler vector.* A first, albeit weak, result is the preceding lemma. A stronger result with a corollary useful for simple estimates is given in the next lemma.

▶ **Lemma 1.5** ([23]). *Let λ_2 be the second-smallest eigenvalue of the unweighted Laplacian of a graph G. Let $\operatorname{diam}(G)$ be the diameter of graph G. Then*

$$\operatorname{diam}(G) \geq \frac{4}{n \cdot \lambda_2}.$$

▶ **Corollary 1.6.** *Using $\operatorname{diam}(G) \leq n$, we get $\lambda_2 \geq \frac{4}{n^2}$*

▶ **Lemma 1.7.** *This is another useful result by Fiedler [16]. Let λ_2 be the second-smallest eigenvalue of $L(G)$. Then,*

$$\lambda_2 \leq \frac{n}{n-1} \cdot \min\{\deg(i), i \in V\}.$$

For Δ the maximum degree of graph G, it immediately follows

$$\lambda_2 \leq \frac{n}{n-1} \cdot \Delta.$$

A stronger relationship between λ_2 and the network's connectivity properties is provided via the graph's Cheeger constant.

▶ **Definition 1.8.** Let $G = (V, E)$ be a graph and $S \subset V$ a subset of the nodes. The *boundary* δS of S is defined as the set of edges having exactly one endpoint in S, i.e.,

$$\delta S = \{(i, j) \in E \mid i \in S, j \in V \setminus S\}.$$

▶ **Definition 1.9.** Let $G = (V, E)$ be a graph. The *isoperimetric number* $i(G)$ of G is defined as

$$i(G) = \min_{\substack{S \subset V \\ |S| \leq |V|/2}} \frac{|\delta S|}{|S|}.$$

It is also called *Cheeger constant* of the graph.

The isoperimetric number of a graph is a measure of how well any subset of the graph is connected to the rest of the graph. Graphs with a high Cheeger constant are also called *expanders*. The following was proven by Mohar.

▶ **Lemma 1.10** ([22]). *Let λ_2 be the second-smallest eigenvalue of $L(G)$, and let $i(G)$ be the isoperimetric number of G. Then,*

$$\frac{i^2(G)}{2\Delta} \leq \lambda_2 \leq 2i(G).$$

This concludes our introduction to spectral graph theory, which suffices for the analysis of identical machines. For machines with speeds, it turns out that a generalized Laplacian is a more expressive quantity.

A.2 Generalized Laplacian Analysis

Recall the speed-matrix S from the introduction. Instead of analyzing the Laplacian L, we are now interested in the *generalized Laplacian*, defined as LS^{-1}. This definition is also used by Elsässer in [11] in the analysis of continuous diffusive load balancing in heterogeneous networks. In this reference, the authors prove a variety of results for the generalized Laplacian, which we restate here in a slightly different language.

It turns out that in the discussion of the properties of this generalized Laplacian, many results carry over from the analysis of the normal Laplacian. The similarity is made manifest by the introduction of a *generalized dot-product*.

▶ **Definition 1.11.** For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we define the *generalized dot-product* with respect to S as

$$\langle \mathbf{x}, \mathbf{y} \rangle_S := \mathbf{x}^T S^{-1} \mathbf{y} = \sum_{i \in V} \frac{x_i \cdot y_i}{s_i}$$

▶ **Lemma 1.12.** *The vector space* $\mathbb{R}^n$ *together with* $\langle\cdot,\cdot\rangle_S$ *forms an* inner product space. *This means that*

(1) $\langle\mathbf{x},\mathbf{y}\rangle_S = \langle\mathbf{y},\mathbf{x}\rangle_S$*, i.e.,* $\langle\cdot,\cdot\rangle_S$ *is symmetric,*

(2) $\langle a\mathbf{x}_1 + b\mathbf{x}_2, y\rangle_S = a\langle\mathbf{x}_1,\mathbf{y}\rangle_S + b\langle\mathbf{x}_2,\mathbf{y}\rangle_S$ *for any scalars* a *and* b*, i.e.,* $\langle\cdot,\cdot\rangle_S$ *is linear in its first argument,*

(3) $\langle\mathbf{x},\mathbf{x}\rangle_S \geq 0$*, with equality if and only if* $\mathbf{x} = 0$*, i.e.,* $\langle\cdot,\cdot\rangle_S$ *is positive definite.*

Proof. All three properties follow immediately from Definition 1.11, provided the s_i are positive, which is true in our case. ◀

▶ Remark. The fact that $\langle\cdot,\cdot\rangle_S$ is an *inner product* allows us to directly apply many results of linear algebra to it. For example, all inner products satisfy the Cauchy-Schwarz inequality, i.e.,

$$\langle\mathbf{x},\mathbf{y}\rangle_S^2 \leq \langle\mathbf{x},\mathbf{x}\rangle_S \cdot \langle\mathbf{y},\mathbf{y}\rangle_S.$$

A proof of this important inequality can be found in every introductory book on Linear Algebra.

Another concept is that of orthogonality. Two vectors $\mathbf{x}$ and $\mathbf{y}$ are called orthogonal to each other, $\mathbf{x}\perp\mathbf{y}$, if $\mathbf{x}\cdot\mathbf{y} = 0$. Analogously, we call $\mathbf{x}$ and $\mathbf{y}$ orthogonal *with respect to* S if $\langle\mathbf{x},\mathbf{y}\rangle_S = 0$.

Let us now collect some of the properties of LS^{-1}. These properties have also been used in [11]. We restate them here using the notation of the generalized dot product.

▶ **Lemma 1.13.** *(Compare Lemma 1 in [11]) Let* L *be the Laplacian of a graph, and let* S *be the speed-matrix,* $S = \operatorname{diag}(s_1, \cdots, s_n)$*. Then the following holds true for the generalized Laplacian* LS^{-1}*.*

(1) The speed-vector $\mathbf{s} = (s_1, \ldots, s_n)^\top$ *is (right-)eigenvector to* LS^{-1} *with eigenvalue* 0*.*

(2) LS^{-1} *is not symmetric any more. It is, however, still positive semi-definite.*

(3) Since LS^{-1} *is not symmetric, we have to distinguish left- and right-eigenvectors. Similar to the spectral theorem of linear algebra, we can find a basis of right-eigenvectors of* LS^{-1} *that are orthogonal with respect to* S*.*

Proof. (1)

$$LS^{-1}\mathbf{s} = L\mathbf{1} = 0$$

via Lemma 1.4. For (2) and (3), suppose that $\mathbf{x}$ is a right-eigenvector of LS^{-1} with eigenvalue λ. If we define $\mathbf{y} := S^{-1/2}\mathbf{x}$, then we have

$$\begin{aligned} && LS^{-1}\mathbf{x} &= \lambda\mathbf{x} \\ \Leftrightarrow && LS^{-1/2}\mathbf{y} &= \lambda S^{1/2}\mathbf{y} \\ \Leftrightarrow && S^{-1/2}LS^{-1/2}\mathbf{y} &= \lambda\mathbf{y}. \end{aligned}$$

This proves that $\mathbf{x}$ is right-eigenvector to LS^{-1} with eigenvalue λ if and only if $S^{-1/2}\mathbf{x}$ is eigenvector to $S^{-1/2}LS^{-1/2}$ with eigenvalue λ. The latter matrix is positive definite, because for every vector $\mathbf{x}$, we have

$$\mathbf{x}^\top S^{-1/2}LS^{-1/2}\mathbf{x} = (S^{-1/2}\mathbf{x})^\top L(S^{-1/2}\mathbf{x}) \geq 0$$

since L itself is positive semi-definite. Now, since $S^{-1/2}LS^{-1/2}$ is symmetric positive semi-definite, all its eigenvalues are real and non-negative and it possesses an orthogonal eigenbasis.

Let us denote the n vectors of the eigenbasis with $\mathbf{y}^k$, $k = 1 \dots n$. As we have just shown, this implies that the vectors $\mathbf{x}^k = S^{1/2}\mathbf{y}^k$ are right-eigenvectors to LS^{-1}. Since $S^{1/2}$ is a matrix of full rank, the $\mathbf{x}^k$ form a basis as well. Their orthogonality with respect to S follows from

$$\begin{aligned}\langle \mathbf{x}^k, \mathbf{x}^l \rangle_S &= (\mathbf{x}^k)^\top S^{-1}\mathbf{x}^l \\ &= (S^{-1/2}\mathbf{x}^k)^\top S^{-1/2}\mathbf{x}^l \\ &= (\mathbf{y}^k)^\top \cdot \mathbf{y}^l = 0 \quad \text{for } k \neq l.\end{aligned}$$

◀

For arbitrary vectors, we know that $\langle \mathbf{x}, LS^{-1}\mathbf{x}\rangle_S \geq 0$ since $S^{-1}LS^{-1}$ is positive semi-definite. The next lemma bounds the generalized dot product of certain vectors with the Laplacian with the second smallest right-eigenvector of it. A similar version can also be found in [11, Section 3].

▶ **Lemma 1.14.** *Let λ_2 denote the second-smallest right-eigenvalue of the generalized Laplacian, LS^{-1}. Let $\mathbf{e}$ be a vector that is orthogonal to the speed vector with respect to S, i.e. $\langle \mathbf{e}, \mathbf{s}\rangle_S = 0$. Then*

$$\langle \mathbf{e}, LS^{-1}\mathbf{e}\rangle_S \geq \lambda_2 \langle \mathbf{e}, \mathbf{e}\rangle_S.$$

Proof. Let $(\lambda_k, \mathbf{v}_k)$ denote the k-smallest eigenvalue and corresponding eigenvector of LS^{-1}. For the speed vector, $\mathbf{s}$, we have $LS^{-1}\mathbf{s} = 0$ (Lemma 1.13). Thus, we can just identify $\mathbf{v}_1 = \mathbf{s}$. Recall from Lemma 1.13 that the $\mathbf{v}_k$ form a basis of $\mathbb{R}^n$. Therefore, $\mathbf{e}$ can be written as a linear combination of these eigenvectors. For some real-valued coefficients β_k, we have

$$\mathbf{e} = \sum_{k=1}^{n} \beta_k \mathbf{v}_k,$$

Since the basis vectors are mutually orthogonal with respect to S, and since $\mathbf{e}$ is orthogonal to $\mathbf{s}$ with respect to S, $\mathbf{s} = \mathbf{v}_1$ does not contribute to the linear combination of $\mathbf{e}$, because

$$0 = \langle \mathbf{e}, \mathbf{v}_1\rangle_S = \beta_1 \langle \mathbf{v}_1, \mathbf{v}_1\rangle.$$

This can only be satisfied if either all speeds are zero or if $\beta_1 = 0$. Therefore, we can write

$$\mathbf{e} = \sum_{k=2}^{n} \beta_k \mathbf{v}_k.$$

Substituting this decomposition into the Bound of Lemma 1.14 yields

$$\begin{aligned}\langle \mathbf{e}, LS^{-1}\sum_{k=2}^{n} \beta_k \mathbf{v}_k\rangle_S &= \sum_{k=2}^{n} \lambda_k \cdot \beta_k \langle \mathbf{e}, \mathbf{v}_k\rangle_S \\ &= \sum_{k=2}^{n} \lambda_k \cdot \beta_k^2 \langle \mathbf{v}_k, \mathbf{v}_k\rangle_S \\ &\geq \lambda_2 \cdot \sum_{k=2}^{n} \langle \beta_k \mathbf{v}_k, \beta_k \mathbf{v}_k\rangle_S = \lambda_2 \cdot \langle \mathbf{e}, \mathbf{e}\rangle_S.\end{aligned}$$

◀

The next technical lemma is needed to relate the spectra of L and LS^{-1}. We require this relation because most of the useful results and bounds for λ_2 apply to the normal Laplacian only.

▶ **Lemma 1.15.** *Let μ_i denote the eigenvalues of LS^{-1} in ascending order and let λ_i denote the eigenvalues of L in ascending order. Finally, let s_i denote the speeds in descending order. Then*

$$\mu_{i+j-1} \geq \frac{\lambda_i}{s_j} \qquad 0 \leq i,j \leq n, \quad 0 \leq i+j-1 \leq n \tag{1}$$

$$\mu_{i+j-n} \leq \frac{\lambda_i}{s_j} \qquad 0 \leq i,j \leq n, \quad 0 \leq i+j-n \leq n. \tag{2}$$

Lemma 1.15. The matrices L and S^{-1} are symmetric positive semi-definite. Hence, their square-roots exist and are unique. Let $X = \sqrt{L}$ and $T = \sqrt{S^{-1}}$. The singular values of X are $\sqrt{\mu_i}$ and those of T are $\sqrt{s_i^{-1}}$. In addition, the singular values of XT are $\sqrt{\lambda_i}$.

By Theorem 4.3, there exist symmetric matrices H_1 with eigenvalues $\log\sqrt{\mu_i}$ and H_2 with eigenvalues $\log\sqrt{s_i^{-1}}$, and the eigenvalues of $H_1 + H_2$ are $\log\sqrt{\lambda_i}$. By Theorem 4.2, these satisfy the inequalities

$$\log\sqrt{\lambda_{i+j-1}} \geq \log\sqrt{\mu_i} + \log\sqrt{s_j^{-1}} = \log\sqrt{\mu_i s_j^{-1}}$$

$$\log\sqrt{\lambda_{i+j-n}} \leq \log\sqrt{\mu_i} + \log\sqrt{s_j^{-1}} = \log\sqrt{\mu_i s_j^{-1}}$$

Since both the logarithm and the square-root are monotone functions, the desired result follows immediately. ◀

▶ **Corollary 1.16.** *Let μ_2 denote the second smallest right eigenvalue of LS^{-1} and let λ_2 denote the second smallest eigenvalue of L. Let $s_{\max} = s_1$ be the largest speed and $s_{\min} = s_n$ the smallest speed. Then*

$$\frac{\lambda_2}{s_{\max}} \leq \mu_2 \leq \frac{\lambda_2}{s_{\min}}.$$

Proof. Let $i = 2, j = 1$ in (1) and $i = 2, j = n$ in (2). ◀

B Proofs from Section 3

Lemma 3.6. (1) By definition, $\Psi_0(x)$ and $\Phi_0(x)$ differ by $m^2/\mathcal{S}$. The total number of tasks, m, and the sum of all speeds, $\mathcal{S}$, are constants. Therefore, the difference between $\Psi_0(x)$ and $\Phi_0(x)$ is constant at any time and we have

$$\Delta\Psi_0(X^t|X^{t-1} = x) = \Delta\Phi_0(X^t|X^{t-1} = x) - \underbrace{\Delta\left[\frac{m^2}{\mathcal{S}}\right]}_{=0}.$$

Hence, both the original and the shifted potential have the same potential drop.

(2) This follows immediately from Definition 1.11 of the generalized dot-product. ◀

B.1 Proofs from Section 3.1

Lemma 3.9. For brevity, we omit the argument x from all quantities. Note that we can look at the drop of either Φ_0 or Ψ_0. Substituting the particular forms of f_{ij} and Λ^0_{ij} (Definitions 3.1 and 3.8) into the bound provided by Lemma 4.1 of Lemma 3.3 in [6], we arrive at

$$\mathbf{E}[\Delta\Psi_0(X^{k+1})|X^k = x] \geq \sum_{(i,j)\in\tilde{E}} \left[\frac{\left(2-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} - \frac{(\ell_i-\ell_j)}{\alpha\cdot d_{i,j}}\right]. \tag{$*$}$$

We define subsets $\tilde{E}$, $\tilde{E}_1$ and $\tilde{E}_2$ of E,

$$\begin{aligned}\tilde{E} &= \left\{(i,j)\in E \mid \ell_i-\ell_j \geq \frac{1}{s_j}\right\}\\ \tilde{E}_1 &= \left\{(i,j)\in \tilde{E} \mid \ell_i-\ell_j \geq \frac{1}{s_i}+\frac{1}{s_j}\right\}\\ \tilde{E}_2 &= \tilde{E}\setminus\tilde{E}_1.\end{aligned}$$

Note that $\tilde{E} = \tilde{E}_1 \cup \tilde{E}_2$ and $\tilde{E}_1 \cap \tilde{E}_2 = \emptyset$. Thus, we can split the sum in $(*)$ into a sum over $\tilde{E}_1$ and a sum over $\tilde{E}_2$. We will now bound these sums individually.

Let $(i,j)\in\tilde{E}_1$ be an edge in $\tilde{E}_1$ so that $\ell_i \geq \ell_j$. Then the definition of $\tilde{E}_1$ and the non-negativity of $\ell_i - \ell_j$ allows us to deduce

$$\ell_i-\ell_j \geq \frac{1}{s_i}+\frac{1}{s_j} \Leftrightarrow \frac{1}{\frac{1}{s_i}+\frac{1}{s_j}}\cdot(\ell_i-\ell_j)^2 \geq \ell_i-\ell_j.$$

This allows us to bound

$$\sum_{(i,j)\in\tilde{E}_1} \left[\frac{\left(2-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} - \frac{(\ell_i-\ell_j)}{\alpha\cdot d_{i,j}}\right] \geq \sum_{(i,j)\in\tilde{E}_1} \frac{\left(1-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}. \tag{$*$}$$

Next, we turn to $\tilde{E}_2$ and bound

$$\sum_{(i,j)\in\tilde{E}_2} \left[\frac{\left(2-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} - \frac{(\ell_i-\ell_j)}{\alpha\cdot d_{i,j}}\right].$$

The sum is over two terms, a positive and a negative one. For the first, positive term, we simply bound

$$\sum_{(i,j)\in\tilde{E}_2} \left[\frac{\left(2-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\right] \geq \sum_{(i,j)\in\tilde{E}_2} \left[\frac{\left(1-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\right]. \tag{$**$}$$

For the edges in $\tilde{E}_2$, we have $\ell_i - \ell_j < 1/s_i + 1/s_j$. This allows us to bound the second, negative term via

$$\sum_{(i,j)\in\tilde{E}_2} \frac{\ell_i-\ell_j}{\alpha\cdot d_{i,j}} \leq \frac{1}{\alpha}\cdot\sum_{(i,j)\in\tilde{E}_2}\frac{1}{d_{i,j}}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right) \tag{$***$}$$

Combining $(*)$, $(**)$ and $(***)$ yields

$$\begin{aligned}\mathbf{E}[\Delta\Psi_0(X^{k+1})|X^k = x] &\geq \sum_{(i,j)\in\tilde{E}} \left[\frac{\left(2-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} - \frac{(\ell_i-\ell_j)}{\alpha\cdot d_{i,j}}\right] \\ &\geq \sum_{(i,j)\in\tilde{E}} \frac{\left(1-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} - \sum_{(i,j)\in\tilde{E}_2}\frac{1}{\alpha\cdot d_{ij}}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right). \qquad (\dagger)\end{aligned}$$

In the next step, we rewrite the sum over $\tilde{E}$ in $(\dagger)$ to a sum over all edges E, using $\tilde{E} = E\setminus(E\setminus\tilde{E})$. It generally holds for any terms $X_{(i,j)}$ that

$$\sum_{(i,j)\in\tilde{E}} X_{(i,j)} = \sum_{(i,j)\in E} X_{(i,j)} - \sum_{(i,j)\in E\setminus\tilde{E}} X_{(i,j)}.$$

We will apply this to $(\dagger)$. In the following, we therefore prove an upper bound on the sum over $E\setminus\tilde{E}$. Without loss of generality, let the nodes i and j of an edge be ordered such that $\ell_i \geq \ell_j$. For edges not in $\tilde{E}$, we have, by definition, $0 \leq \ell_i - \ell_j \leq \frac{1}{s_j}$, so this part can be bound by

$$\begin{aligned}\sum_{(i,j)\in E\setminus\tilde{E}} \frac{\left(1-\frac{2}{\alpha}\right)}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot(\ell_i-\ell_j)^2 &\leq \sum_{(i,j)\in E\setminus\tilde{E}} \frac{1}{\alpha\cdot d_{ij}}\cdot\frac{s_i}{s_i+s_j}\cdot\frac{1}{s_j} \\ &\leq \sum_{(i,j)\in E\setminus\tilde{E}} \frac{1}{\alpha\cdot d_{ij}}\cdot\left(1-\frac{s_j}{s_i+s_j}\right)\cdot\frac{1}{s_j} \\ &= \sum_{(i,j)\in E\setminus\tilde{E}} \frac{1}{\alpha\cdot d_{i,j}}\cdot\left(\frac{1}{s_j}-\frac{1}{s_i+s_j}\right) \\ &\leq \sum_{(i,j)\in E\setminus\tilde{E}} \frac{1}{\alpha\cdot d_{i,j}}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right).\end{aligned}$$

This bound has the same form as the bound in $(***)$, only that it goes over $E\setminus\tilde{E}$ instead of $\tilde{E}_2$. These two sets are disjunct, since $\tilde{E}_2 \subset \tilde{E}$. Therefore, we can combine the two sums into a single sum over $\tilde{E}_2 \cup (E\setminus\tilde{E}) = E\setminus\tilde{E}_1$. We then obtain from the following bound from $(\dagger)$.

$$\begin{aligned}\mathbf{E}[\Delta\Psi_0(X^{k+1})|X^k = x] \geq &\sum_{(i,j)\in E}\left[\frac{\left(1-\frac{2}{\alpha}\right)\cdot(\ell_i-\ell_j)^2}{\alpha\cdot d_{i,j}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\right] \\ &- \sum_{(i,j)\in E\setminus\tilde{E}_1}\frac{1}{\alpha\cdot d_{ij}}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right). \qquad (\dagger\dagger)\end{aligned}$$

The first term of this bound already has the desired form. We will now bound the second term. Since it is negative, we have to upper bound the sum itself. First, note that $E\setminus\tilde{E}_1$ is a subset of E. As the term inside the sum is non-negative, we can write

$$\sum_{(i,j)\in E\setminus\tilde{E}_1}\frac{1}{\alpha\cdot d_{ij}}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right) \leq \frac{1}{\alpha}\cdot\sum_{(i,j)\in E}\left[\frac{1}{d_{ij}\cdot s_i}+\frac{1}{d_{ij}\cdot s_j}\right].$$

Recall that d_{ij} is defined as $\max\{\deg(i), \deg(j)\}$, so we can bound

$$\begin{aligned}\frac{1}{\alpha} \cdot \sum_{(i,j)\in E} \left[\frac{1}{d_{ij} \cdot s_i} + \frac{1}{d_{ij} \cdot s_j}\right] &\leq \frac{1}{\alpha} \cdot \sum_{(i,j)\in E} \left[\frac{1}{\deg(i) \cdot s_i} + \frac{1}{\deg(j) \cdot s_j}\right] \\ &= \frac{1}{\alpha} \cdot \sum_{i\in V} \sum_{j\in \mathrm{Adj}(i)} \frac{1}{\deg(i) \cdot s_i} \\ &= \frac{1}{\alpha} \cdot \sum_{i\in V} \frac{1}{s_i} \leq \frac{n}{\alpha}.\end{aligned}$$

Inserting this bound into (††) yields the result. ◀

Lemma 3.10. We start from the bound obtained in Lemma 3.9. In the course of this proof, we will use Lemma 1.14 for the task deviation vector $\mathbf{e}$. In order to use this lemma, we have to show that $\mathbf{e}$ satisfies the lemma's condition, i.e., $\langle \mathbf{e}, \mathbf{s}\rangle_S = 0$. This follows via

$$\langle \mathbf{e}, \mathbf{s}\rangle_S = \sum_{i\in V} \frac{e_i \cdot s_i}{s_i} = \sum_{i\in V} e_i = 0.$$

We can now begin with the main proof.

$$\begin{aligned}\mathbf{E}[\Delta\Psi_0(X^{k+1})|X^k = x] &\geq \sum_{(i,j)\in E} \left[\frac{\left(1 - \frac{2}{\alpha}\right) \cdot (\ell_i(x) - \ell_j(x))^2}{\alpha \cdot d_{i,j} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right)}\right] - \frac{n}{\alpha} \\ &\geq \frac{\frac{1}{2}}{4 \cdot s_{\max} \cdot \Delta \cdot 2} \cdot \sum_{(i,j)\in E} \left(\ell_i - \frac{m}{\mathcal{S}} + \frac{m}{\mathcal{S}} - \ell_j\right)^2 - \frac{n}{4 \cdot s_{\max}} \\ &\geq \frac{1}{16\Delta} \cdot \frac{1}{s_{\max}} \cdot \sum_{(i,j)\in E} \left(\frac{e_i}{s_i} - \frac{e_j}{s_j}\right)^2 - \frac{n}{4 \cdot s_{\max}} \\ \text{(Lem. 1.13)} &= \frac{1}{16\Delta} \cdot \frac{1}{s_{\max}} \cdot \left(S^{-1}\mathbf{e}\right)^\top L \left(S^{-1}\mathbf{e}\right) - \frac{n}{4 \cdot s_{\max}} \\ &= \frac{1}{16\Delta} \cdot \frac{1}{s_{\max}} \cdot \langle \mathbf{e}, LS^{-1}\mathbf{e}\rangle_S - \frac{n}{4 \cdot s_{\max}} \\ \text{(Lem. 1.14, Cor. 1.16)} &\geq \frac{1}{16\Delta} \cdot \frac{1}{s_{\max}} \cdot \frac{\lambda_2}{s_{\max}} \langle \mathbf{e}, \mathbf{e}\rangle_S - \frac{n}{4 \cdot s_{\max}} \\ \text{(Lem. 3.6)} &= \frac{1}{16\Delta} \cdot \frac{1}{s_{\max}} \cdot \frac{\lambda_2}{s_{\max}} \cdot \Psi_0 - \frac{n}{4 \cdot s_{\max}}\end{aligned}$$

◀

Lemma 3.11. First, we find a bound for $\mathbf{E}[\Psi_0(X^t)|X^{t-1} = x]$, i.e., the expected value of the potential if the previous state was x. There, we have

$$\begin{aligned}\mathbf{E}[\Psi_0(X^t)|X^{t-1} = x] &= \Psi_0(x) - \mathbf{E}[\Delta\Psi_0(X^t)|X^{t-1} = x] \\ \text{(Lem. 3.10)} \quad &\leq \left(1 - \frac{2}{\gamma}\right) \cdot \Psi_0(x) + \frac{n}{4 \cdot s_{\max}}.\end{aligned}$$

Using this, we now use the tower property of iterated expectation to obtain

$$\begin{aligned}\mathbf{E}[\Psi_0(X^t)] &= \mathbf{E}[\mathbf{E}[\Psi_0(X^t)|X^{t-1}]] \\ &\leq \left(1 - \frac{2}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^{t-1})] + \frac{n}{4 \cdot s_{\max}}.\end{aligned}$$

◀

Lemma 3.13. We can rearrange the condition $\mathbf{E}[\Psi_0(X^t)] \geq \psi_c$ as follows.

$$\begin{aligned} & \mathbf{E}[\Psi_0(X^t)] \geq \frac{8 \cdot n \cdot \Delta \cdot s_{\max}}{\lambda_2} \\ \Leftrightarrow \quad & \underbrace{\frac{\lambda_2}{32\Delta} \cdot \frac{1}{s_{\max}^2}}_{=1/\gamma} \cdot \mathbf{E}[\Psi_0(X^t)] \geq \frac{n}{4 \cdot s_{\max}}. \end{aligned}$$

We insert this into the bound from Lemma 3.11 to obtain

$$\begin{aligned} \mathbf{E}[\Psi_0(X^{t+1})] &\leq \left(1 - \frac{2}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^t)] + \frac{n}{4 \cdot s_{\max}} \\ &\leq \left(1 - \frac{2}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^t)] + \frac{1}{\gamma}\,\mathbf{E}[\Psi_0(X^t)] \\ &\leq \left(1 - \frac{1}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^t)]. \end{aligned}$$

◀

Lemma 3.14. The proof is by induction on t, with $t = 0$ as the base case. There, we have

$$\mathbf{E}[\Psi_0(X^0)] \leq \left(1 - \frac{1}{\gamma}\right)^0 \cdot \mathbf{E}[\Psi_0(X^0)].$$

Next, let the claim be true for a given t. Either there is a $t' < t$ such that $\mathbf{E}[\Psi_0(X^{t'})] \leq \psi_c$ and we are done. Otherwise, we can apply Lemma 3.13 to obtain.

$$\begin{aligned} \mathbf{E}[\Psi_0(X^{t+1})] &\leq \left(1 - \frac{1}{\gamma}\right) \cdot \mathbf{E}[\Psi_0(X^t)] \\ \text{(Induction Hypothesis)} \quad &\leq \left(1 - \frac{1}{\gamma}\right)^{t+1} \cdot \mathbf{E}[\Psi_0(X^0)]. \end{aligned}$$

◀

Lemma 3.15. (1) We write down the inequality we want to prove and then show that the T given in the lemma makes it true. Since we use Lemma 3.14 to connect the expectation value of Ψ_0 at time T to its value at time $t = 0$, we first note that

$$\Psi_0(X^0) \leq m^2, \tag{$*$}$$

as the largest potential is obtained by the largest imbalance, i.e., assigning all m tasks to the slowest node. Then, we can deduce

$$\begin{aligned} & \mathbf{E}[\Psi_0(X^T)] \leq \psi_c \\ \text{(Lem. 3.14)} \Leftarrow \quad & (1 - 1/\gamma)^T \cdot \Psi_0(X^0) \leq \frac{16\Delta \cdot n \cdot s_{\max}}{\lambda_2} \\ (*) \Leftarrow \quad & T \cdot \ln(1 - 1/\gamma) + \ln\left(m^2\right) \leq \ln\left(\frac{16\Delta \cdot n \cdot s_{\max}}{\lambda_2}\right). \end{aligned}$$

Next we use Lemma 1.7, which states $\lambda_2 \leq n/(n-1) \cdot \Delta$, together with $n - 1 \geq n/2$ for

$n \geq 2$ to further rewrite the bound.

$$\begin{aligned} & T \cdot \ln(1-1/\gamma) + \ln\left(m^2\right) \leq \ln\left(\frac{16\Delta \cdot n \cdot s_{\max}}{\lambda_2}\right) \\ \text{(Lem. 1.7)} \Leftarrow \quad & -\frac{1}{\gamma} \cdot T + 2\ln(m) \leq \ln\left(16(n-1)\right) + \ln(s_{\max}) \\ \Leftarrow \quad & -\frac{1}{\gamma} \cdot T + 2\ln(m) \leq \ln(8n) \\ \Leftarrow \quad & -\frac{1}{\gamma} \cdot T + 2\ln(m) \leq \ln(n) \end{aligned}$$

This can be rearranged to yield the condition on T,

$$T \geq 2\gamma \cdot \ln\left(\frac{m}{n}\right).$$

(2) To prove a lower bound on the probability that $\Psi_0(X^t) \leq 4 \cdot \psi_c$, we prove an upper bound on the complementary event. Let Y denote the random variable $\Psi_0(X^t)$, with the t from part (1) of this lemma, i.e. the t for which $\mathbf{E}[\Psi_0(X^t)] \leq \psi_c$. Then, we have

$$\mathbf{Pr}[\Psi_0(X^t) > 4 \cdot \psi_c] \leq \mathbf{Pr}\left[\Psi_0(X^t) > 4 \cdot \mathbf{E}[\Psi_0(X^t)]\right].$$

Applying Markov's inequality to this result immediately yields

$$\mathbf{Pr}[\Psi_0(X^t) > 4 \cdot \mathbf{E}[\Psi_0(X^t)]] \leq \frac{1}{4}.$$

Hence,

$$\mathbf{Pr}[\Psi_0(X^t) \leq 4 \cdot \psi_c] \geq 1 - \frac{1}{4} = \frac{3}{4}.$$

◀

Observation 3.16. We omit the argument x for brevity. Let us begin with the first inequality. Let k be the index of the node for which $|e_k/s_k|$ is maximized. Then

$$L_\Delta^2 = \frac{e_k^2}{s_k^2} \leq \frac{e_k^2}{s_k^2} \cdot s_k \leq \sum_{i \in V} \frac{e_i^2}{s_i} = \Psi_0.$$

The second inequality follows from

$$\Psi_0 = \sum_{i \in V} \frac{e_i^2}{s_i} = \sum_{i \in V} \frac{e_i^2}{s_i^2} \cdot s_i \leq L_\Delta^2 \sum_{i \in V} s_i = L_\Delta^2 \cdot \mathcal{S}$$

◀

Lemma 3.17. From Observation 3.16 we have $L_\Delta^2(x) \leq \Psi_0(x)$. Hence, we have for all states with $\Psi_0(x) \leq 4 \cdot \psi_c$:

$$L_\Delta \leq 16 \cdot \sqrt{\frac{n \cdot \Delta}{\lambda_2} \cdot s_{\max}} \leq 8 \cdot n^2 \cdot s_{\max} =: a$$

where Corollary 1.6 states that $1/\lambda_2 \leq n^2/4$.

Next, we define $\varepsilon = 2/(1+\delta)$. The condition for an ε-approximate Nash equilibrium is that for every edge $(i,j) \in E$ we have

$$(1-\varepsilon) \cdot \frac{w_i}{s_i} \leq \frac{w_j + 1}{s_j}.$$

To prove that this is the case, note that the definition of L_Δ ensures that

$$\left| \frac{w_i}{s_i} - \frac{m}{\mathcal{S}} \right| \leq L_\Delta \Longrightarrow \frac{w_i}{s_i} \leq L_\Delta + \frac{m}{\mathcal{S}} \leq a + \frac{m}{\mathcal{S}}$$

and, analogously

$$\frac{w_j}{s_j} \geq \frac{m}{\mathcal{S}} - a.$$

With this, we have

$$\begin{aligned} & & (1-\varepsilon) \cdot \frac{w_i}{s_i} &\leq \frac{w_j+1}{s_j} \\ &\Leftarrow & \frac{\delta - 1}{\delta + 1} \cdot \left(a + \frac{m}{\mathcal{S}} \right) &\leq \left(\frac{m}{\mathcal{S}} - a \right) \\ &\Leftarrow & \delta \cdot a - \frac{m}{\mathcal{S}} &\leq \frac{m}{\mathcal{S}} - \delta \cdot a \\ &\Leftrightarrow & m &\geq \delta \cdot a \cdot \mathcal{S} = 8 \cdot \delta \cdot s_{\max} \cdot \mathcal{S} \cdot n^2. \end{aligned}$$

Thus, if the number of tasks is sufficiently high, we have an ε-approximate Nash equilibrium. ◀

B.2 Proofs from Section 3.2

Observation 3.20. For brevity, we omit the argument x. We begin with (1). This is simple algebra.

$$\begin{aligned} \Psi_1 &= \Phi_1 - \frac{m^2}{\mathcal{S}} - \frac{m \cdot n}{\mathcal{S}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a} \right) \\ &= \sum_{i \in V} \frac{w_i \cdot (w_i + 1)}{s_i} - \frac{m^2}{\mathcal{S}} - \frac{m \cdot n}{\mathcal{S}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a} \right) \\ &= \sum_{i \in V} \left[\frac{(e_i + m/\mathcal{S} \cdot s_i)^2}{s_i} + \frac{e_i + m/\mathcal{S} \cdot s_i}{s_i} \right] - \frac{m^2}{\mathcal{S}} - \frac{m \cdot n}{\mathcal{S}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a} \right) \\ &= \sum_{i \in V} \left[\frac{e_i^2 + e_i}{s_i} \right] + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a} \right) \\ &= \sum_{i \in V} \frac{\left(e_i + \frac{1}{2}\right)^2}{s_i} - \frac{n}{4 \cdot \bar{s}_a}. \end{aligned}$$

Next, we prove (2). From the form in (1) it might appear that $\Psi_1(x)$ can be negative. Note, however, that the task deviation vector $\mathbf{e}$ satisfies $\sum_i e_i = 0$. We can use the technique of Lagrange multiplicators to find the minimum of $\Psi_1(x)$ under the constraint that the deviations sum to 0. For an additional parameter λ, the so called *Lagrange multiplier*, we define the Lagrange function

$$\mathcal{L}(e_1, \ldots, e_n; \lambda) = \sum_{i \in V} \frac{\left(e_i + \frac{1}{2}\right)^2}{s_i} - \lambda \cdot \sum_{i \in V} e_i.$$

The constrained minimum of $\Psi_1(x)$ is obtained for the solution of

$$\frac{\partial \mathcal{L}}{\partial e_i} = 0, \qquad \frac{\partial \mathcal{L}}{\partial \lambda} = 0.$$

Carrying out the calculation shows that, indeed, minimum value of Ψ_1 therefore is 0.

(3) is obtained by first rewriting Φ_1 as

$$\Phi_1 = \sum_{i \in V} \frac{w_i \cdot (w_i + 1)}{s_i} = \Phi_0 + \sum_{i \in V} \frac{\ell_i}{s_i} = \Phi_0 + \sum_{i \in V} \frac{e_i}{s_i} + \frac{m \cdot n}{\mathcal{S}}.$$

Recall the definition of $\Psi_0 = \Phi_0 - m^2/\mathcal{S}$. Hence

$$\begin{aligned}\Psi_1 &= \Phi_1 - \frac{m^2}{\mathcal{S}} - \frac{m \cdot n}{\mathcal{S}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right) \\ &= \Phi_0 + \sum_{i \in V} \frac{e_i}{s_i} - \frac{m^2}{\mathcal{S}} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right) \\ &= \Psi_0 + \sum_{i \in V} \frac{e_i}{s_i} + \frac{n}{4} \cdot \left(\frac{1}{\bar{s}_h} - \frac{1}{\bar{s}_a}\right).\end{aligned}$$

(4) follows from the definition of $\Psi_1(x)$ by observing that apart from the term $\Phi_1(x)$, everything else is constant. ◀

Lemma 3.21. We have

$$\begin{aligned} & & \ell_i - \ell_j &> \frac{1}{s_j} \\ &\Leftrightarrow & w_i \cdot s_j - w_j \cdot s_i &> s_i \\ &\Leftrightarrow & w_i \cdot n_j \cdot \epsilon - w_j \cdot n_i \cdot \epsilon &> n_i \cdot \epsilon \\ &\Leftrightarrow & w_i \cdot n_j - w_j \cdot n_i &> n_i \end{aligned}$$

Since the lefthand side and righthand side expressions are integers, this leads to

$$\begin{aligned} &\Leftrightarrow & w_i \cdot n_j - w_j \cdot n_i &\geq n_i + 1 \\ &\Leftrightarrow & w_i \cdot s_j - w_j \cdot s_i &\geq s_i + \epsilon \\ &\Leftrightarrow & \ell_i - \ell_j &\geq \frac{1}{s_j} + \frac{\epsilon}{s_i \cdot s_j} \end{aligned}$$

◀

Lemma 3.22. Apply 3.21 to 4.1, then bound the result.

$$\begin{aligned}
\mathbf{E}[\Delta\Psi_1] &\geq \sum_{(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\left[\left(2-\frac{2}{\alpha}\right)\cdot\left(\frac{1}{s_j}+\frac{\epsilon}{s_i\cdot s_j}\right)-\frac{2}{s_j}\right]\\
&= \sum_{(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\left[\frac{2\epsilon}{s_i\cdot s_j}-\frac{2\epsilon}{\alpha\cdot s_i\cdot s_j}-\frac{2}{\alpha\cdot s_j}\right]\\
&\geq \sum_{(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\left[\frac{2\epsilon}{s_i\cdot s_j}-\frac{2\epsilon^2}{4\cdot s_{\max}\cdot s_i\cdot s_j}-\frac{2\epsilon}{4\cdot s_{\max}\cdot s_j}\right]\\
&\geq \sum_{(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\left[\frac{2\epsilon}{s_i\cdot s_j}-\frac{\epsilon}{2\cdot s_i\cdot s_j}-\frac{\epsilon}{2\cdot s_i\cdot s_j}\right]\\
&= \sum_{(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)}\cdot\frac{\epsilon}{s_i\cdot s_j}\\
&\geq \sum_{(i,j)\in\tilde{E}} \frac{\epsilon}{4\cdot s_{\max}\cdot d_{ij}\cdot 2\cdot s_{\max}}\cdot\frac{1}{s_j}\\
&\geq \sum_{(i,j)\in\tilde{E}} \frac{\epsilon^2}{8\cdot\Delta s_{\max}^3} = |\tilde{E}|\cdot\frac{\epsilon^2}{8\Delta\cdot s_{\max}^3} \geq \frac{\epsilon^2}{8\Delta\cdot s_{\max}^3}.
\end{aligned}$$

◀

Lemma 3.23. Observation 3.20 (3) states that with $\mathbf{e}$ denoting the task deviation vector and $\mathbf{1}$ denoting the vector $(1,\dots,1)^\top$, we have

$$\Psi_1(x) = \Psi_0(x) + \langle\mathbf{e},\mathbf{1}\rangle_S + \frac{n}{4}\cdot\left(\frac{1}{\bar{s}_h}-\frac{1}{\bar{s}_a}\right).$$

It remains to bound the dot-product in the equation above. Since $\langle\cdot,\cdot\rangle_S$ is an inner product, it obeys the Cauchy-Schwarz inequality and we have

$$|\langle\mathbf{e},\mathbf{1}\rangle_S|^2 \leq \langle\mathbf{e},\mathbf{e}\rangle_S\cdot\langle\mathbf{1},\mathbf{1}\rangle_S = \Psi_0(x)\cdot\sum_{i\in V}\frac{1}{s_i} = \Psi_0(x)\cdot\frac{n}{\bar{s}_h}.$$

◀

Lemma 3.24. The first part is obtained from simply inserting the definition. Since for times $t \leq T$ the system is not in a Nash equilibrium, we can use Corollary 3.22 to write

$$\begin{aligned}
\mathbf{E}[Z_t|Z_{t-1}=z] &= \mathbf{E}[\Psi_1(X^t)|\Psi_1(X^{t-1})+(t-1)V=z]\\
&\leq \Psi_1(X^{t-1}) - V + tV = z-(t-1)V-V+tV = z.
\end{aligned}$$

For the second part, note that

$$\mathbf{E}[Z_t] = \mathbf{E}[\mathbf{E}[Z_t|Z_{t-1}=z]] \leq \mathbf{E}[Z_{t-1}].$$

◀

Corollary 3.25. For time steps $t \leq T$, we have $t\wedge T = t$ and we can just use Lemma 3.24. Note that the expected constant drop in potential Ψ_1 only depends on the *current* value of the potential. Hence

$$\mathbf{E}[Z_t|Z_0,Z_1,\dots Z_{t-1}] = \mathbf{E}[Z_{t\wedge T}|Z_{t-1}=z] \leq z.$$

For time steps $t > T$, we have

$$\mathbf{E}[Z_t | Z_0, Z_1, \dots Z_{t-1}] = \mathbf{E}[Z_T | Z_0, \dots Z_{T-1}] = \mathbf{E}[Z_T | Z_{T-1} = z] \le z.$$

◀

Corollary 3.26. First, note that the random variable T is a *stopping time* for $Z_{t \wedge T}$, because the event that $T = t$ for some time step t depends only on the state X^t and, most importantly, does not depend on some X^t for $t > T$. The Optional Stopping Theorem allows us to obtain the claim of the corollary for a stopping time T if $\mathbf{E}[T] < \infty$ and $\mathbf{E}[|Z_{t+1\wedge T} - Z_{t\wedge T}|] < c$ for some constant c. The first condition follows from the constant drop in the potential Ψ_1 as long as the system is not in a Nash equilibrium. The second condition follows from

$$\mathbf{E}[|Z_{t+1\wedge T} - Z_{t\wedge T}|] \le |\Delta\Psi_1(X^{t+1})| \le \max_x \Psi_1(x) + V.$$

For any given system, this expression is clearly a constant. The Optional Stopping Theorem for super-martingales now states that

$$\mathbf{E}[Z_T] = \mathbf{E}[Z_{T\wedge T}] \le \mathbf{E}[Z_0] = \Psi_1(X^0).$$

◀

C Proofs from Section 4

Lemma 4.3. As in the original reference, we introduce random variables A_i and C_i for the tasks *abandoning* and *coming to* node i, but now they count the *weight* of these tasks, not only their number. For the C_i, we again split it into Z_{ji} where Z_{ji} counts the weight migrating from j to i. Then

$$\mathbf{Var}[C_i] = \sum_{j:(j,i)\in\tilde{E}} \mathbf{Var}[Z_{ji}].$$

$$\mathbf{Var}[Z_{ji}] = \sum_{\ell\in x_j} \mathbf{Var}[w_\ell^{ji}].$$

Here w_ℓ^{ji} is the random variable that is w_ℓ if task ℓ moves from j to i and 0 otherwise. This variable follows a *Bernoulli distribution*. If p is the probability for the event to occur and if x is the value of the event, then the variance is

$$\mathbf{Var}[\mathsf{Ber}(x,p)] = x^2 \cdot p \cdot (1-p) \le x^2 p.$$

This allows us to write

$$\begin{aligned}
\mathbf{Var}[Z_{ji}] &= \sum_{l\in x_j} \mathbf{Var}[w_\ell^{ji}] \\
&\le \sum_{l\in x_j} w_\ell^2 \cdot \frac{\ell_j - \ell_i}{\alpha \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot W_j} \\
&\le \sum_{l\in x_j} w_\ell \cdot \frac{\ell_j - \ell_i}{\alpha \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right) \cdot W_j} \\
&= \frac{\ell_j - \ell_i}{\alpha \cdot d_{ij} \cdot \left(\frac{1}{s_i} + \frac{1}{s_j}\right)} = f_{ij},
\end{aligned}$$

where we use that $w_\ell^2 \le w_\ell$ since all tasks have weight at most 1. Hence

$$\mathbf{Var}[C_i] = \sum_{j:(j,i)\in\tilde{E}} f_{ij}.$$

Similarly, we define the random variable $A_{\ell i}$ that is w_l if task ℓ abandons node i and 0 otherwise. It is also Bernoulli-distributed.

$$\begin{aligned}
\mathbf{Var}[A_i] &= \sum_{\ell\in x_i} \mathbf{Var}[A_{\ell i}] \\
&= \sum_{\ell\in x_i} \mathbf{Var}\left[\mathsf{Ber}\left(w_\ell; \sum_{j:(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W_i}\right)\right] \le \sum_{\ell\in x_i} w_\ell^2 \cdot \sum_{j:(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W_i} \\
&\le \sum_{j:(i,j)\in\tilde{E}} \sum_{\ell\in x_i} w_\ell \cdot \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)\cdot W_i} = \sum_{j:(i,j)\in\tilde{E}} \frac{\ell_i-\ell_j}{\alpha\cdot d_{ij}\cdot\left(\frac{1}{s_i}+\frac{1}{s_j}\right)} = \sum_{j:(i,j)\in\tilde{E}} f_{ij}.
\end{aligned}$$

When we add the variance of C_i and A_i and sum over all nodes, we get, in formal analogy to the unweighted case,

$$\sum_i \frac{\mathbf{Var}[W_i(X^t)|X^{t-1}=x]}{s_i} = \sum_{ij} f_{ij}\left(\frac{1}{s_i}+\frac{1}{s_j}\right)$$

◀

D Auxiliary Results

In this part we collect results from other papers that are essential in our own proofs.

▶ **Lemma 4.1** (Lemma 3.3 in [6]). *For any step t and any state x,*

$$\mathbf{E}[\Delta\Phi_r(X^t)|X^{t-1}=x] \ge \sum_{(i,j)\in\tilde{E}(x)} f_{ij}(x)\cdot\left(\Lambda_{ij}^r(x) - \frac{1}{s_i} - \frac{1}{s_j}\right).$$

▶ **Theorem 4.2** ([28], [7]). *Let A, B and $C = A+B$ be Hermitian matrices with eigenvalues α_i, β_i and γ_i in ascending order. Then*

$$\begin{aligned}
\gamma_{i+j-1} &\ge \alpha_i + \beta_j \\
\gamma_{i+j-n} &\le \alpha_i + \beta_j
\end{aligned}$$

▶ **Theorem 4.3** (Theorem B from [21]). *The following conditions are equivalent*

1. *There exist matrices $A_i \in \mathrm{GL}(n,\mathbb{R})$ with given singular spectra*

$$\sigma_i = \sigma(A_i) \quad \textit{and} \quad \tilde{\sigma} = \sigma(A_1A_2\cdots A_N).$$

2. *There exist symmetric $n\times n$-matrices H_i with spectra*

$$\lambda(H_i) = \log\sigma_i \quad \textit{and} \quad \lambda(H_1+H_2+\cdots+H_N) = \log\sigma,$$

that is, the eigenvalues of H_i are the logarithms of the singular values of A_i, and the eigenvalues of $\sum_i H_i$ are the logarithms of the singular values of $\prod_i A_i$.